Design and Optimization of Low-thrust Orbit Transfers

Seungwon Lee, Paul von Allmen, Wolfgang Fink, Anastassios E. Petropoulos, and Richard J. Terrile
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA
818-393-7720
Seungwon.Lee@jpl.nasa.gov

Abstract— NASA's space missions Dawn and JIMO will use low-thrust propulsion for multi-revolution orbit transfers around a central body. Here we address the problem of designing low-thrust orbit transfers between arbitrary orbits in an inverse-square gravity field by using evolutionary algorithms to drive parameter selection in a Lyapunov feedback control law (the Q-law). We develop an efficient and efficacious method to assess, with reasonable accuracy, the trade off between propellant mass and flight time (i.e., to find the Pareto front for these two quantities), and to provide the time history of the state variables and the thrust vector for any chosen point on the Pareto front.

TABLE OF CONTENTS

1 INTRODUCTION
2 Q-LAW
3 Q-LAW OPTIMIZATION WITH EVOLUTIONARY ALGORITHMS
4 ORBIT-TRANSFER RESULTS
5 CONCLUSIONS
6 ACKNOWLEDGMENTS
7 APPENDIX

1. INTRODUCTION

NASA's future space missions Dawn and JIMO will use electric propulsion for inter-planetary cruise and orbital operations. The strength of electric propulsion is that despite its low thrust levels, the momentum transfer to the spacecraft per kilogram of expelled propellant is ten or twenty times greater than for chemical propulsion. However, the control of continually-thrusting, low-thrust spacecraft poses a challenging design problem, particularly for orbit transfers around a central body. Third-body perturbations and perturbations from non-spherical mass distributions in the central body often dominate the thrust. Yet even without these perturbations, where only a simple inverse-square gravity field is considered, low-thrust orbit transfers are particularly challenging to design due to the large number of revolutions around the central body and the difficulty of selecting thrust directions and thrust arc locations. Such transfers have been studied at least since the 1950s [1][2]. While the ultimate goal may be to include the full gravity field, as a first step we address in this paper the problem of designing low-thrust orbit transfers between arbitrary initial and final orbits in an inverse-square gravity field. We couple the Q-law, which is a Lyapunov feedback control law developed by Petropoulos [3][4], with evolutionary algorithms to select parameters in the Q-law.

It has been demonstrated that the Q-law, with a reasonable set of control parameters, efficiently finds approximate Pareto-optimal solutions (i.e., a propellant-optimal solution for a given flight time or a flight-time-optimal solution for a given propellant requirement) [3][4]. On the other hand, a grid sampling of the Q-law parameters suggests that a better solution can be found if optimized Q-law parameters are used [4]. Finding an optimal set of the Q-law parameters for all possible orbit transfers is analytically impossible and can be computationally expensive without good heuristic algorithms. There is no guarantee that a single set of the Q-law parameters is superior for all types of orbit transfers. It is also not expected that one particular set is superior for all propellant requirements and flight times of a specific orbit transfer. In this paper, we demonstrate that a genetic algorithm and a simulated-annealing-related algorithm efficiently optimize the Q-law parameters, and thus improve the estimation of the propellant-mass and flight-time requirements for various orbit transfers. The five orbit transfers considered are those presented by Petropoulos [4], or slight variants thereof, and are compared with optimal solutions from the literature, where these are available.

2. Q-LAW

The Q-law was developed by Petropoulos [3][4] in order to provide good initial guesses for propellant-optimal low-thrust orbit transfers. The Q-law determines when and at what angles to thrust based on the proximity quotient termed Q. The function Q judiciously quantifies the proximity of the osculating orbit to the target orbit. In the Q-law, the central body is modeled as a point mass, and no perturbing forces are considered. We summarize in the remainder of this section the Q-law from Ref. [4].

The Q-law consists of two main control rules. 1) The Q-law chooses the thrust angles which reduce Q most quickly at the current instant. 2) The Q-law determines whether to thrust or coast according to a given thrust effectivity threshold $\eta_{\rm cut} \in$

0-7803-8870-4/05/$20.00/
IEEEAC paper # 2.1204

$[0, 1]$ as follows:

$$\text{thrust,} \qquad \text{if } \frac{\min_{\alpha,\beta} \dot{Q}}{\min_{\alpha,\beta,\theta} \dot{Q}} \geq \eta_{\text{cut}} \tag{1}$$

$$\text{coast,} \qquad \text{if } \frac{\min_{\alpha,\beta} \dot{Q}}{\min_{\alpha,\beta,\theta} \dot{Q}} < \eta_{\text{cut}}, \tag{2}$$

where $\dot{Q}$ is the time rate of change of Q, α and β are the thrust angles (more specifically, the azimuthal and elevation angles of the thrust with the pole being given by the osculating orbital angular momentum), and θ is the true anomaly of the osculating orbit. $\min_{\alpha,\beta} \dot{Q}$ is the minimum of $\dot{Q}$ over α and β at a given θ, whereas $\min_{\alpha,\beta,\theta} \dot{Q}$ is the minimum of $\dot{Q}$ over α, β, and θ. Thus, η_{cut} is a handle to control the minimum tolerated effectivity of the thrust. In general, a larger η_{cut} leads to a smaller propellant mass used and a longer flight time.

In some situations, thrust on-off jitter occurs if Eqs. (1) and (2) are followed strictly. These situations arise at true anomalies where thrusting would reduce the effectivity, while coasting would increase it. The jitter is eliminated by imposing a minimum thrust-arc length, typically just a few degrees.

Orbit propagation is done by numerically integrating Gauss's form of the variational equations for the orbit elements [13] using a 5^{th}-6^{th}-order Runge-Kutta-Fehlberg algorithm. At each integration step, the Q-law provides an indication of whether to apply thrust or not, and, if so, in which direction. The thrust, when on, and the specific impulse are assumed constant.

The proximity quotient Q, which serves as a candidate Lyapunov function in the Q-law, is defined as follows:

$$Q = (1 + W_P P) \sum_{\boldsymbol{\varpi}} W_{\boldsymbol{\varpi}} S_{\boldsymbol{\varpi}} \left[\frac{d(\boldsymbol{\varpi}, \boldsymbol{\varpi}_T)}{\dot{\boldsymbol{\varpi}}_{\text{xx}}} \right]^2 \qquad \text{for } \boldsymbol{\varpi} = a, e, i, \omega, \Omega. \tag{3}$$

The five orbital elements ($\boldsymbol{\varpi}$) are the semimajor axis (a), eccentricity (e), inclination (i), argument of periapsis (ω), and longitude of the ascending node (Ω); W_P and the $W_{\boldsymbol{\varpi}}$ are scalar weights greater than or equal to zero; the subscript T denotes the target orbit element value (without subscript, the osculating value is indicated); $\dot{\boldsymbol{\varpi}}_{\text{xx}}$ denotes the maximum over thrust direction and over true anomaly on the osculating orbit of the rate of change of the orbit element (due to thrust). The analytical expressions for $\dot{\boldsymbol{\varpi}}_{\text{xx}}$ are available in Ref. [3]; P is a penalty function; $S_{\boldsymbol{\varpi}}$ is a scaling function; and $d(\boldsymbol{\varpi}, \boldsymbol{\varpi}_T)$ is a distance function. The penalty function is used in the present paper to enforce minimum-periapsis-radius constraints and takes the form

$$P = \exp\left[k \left(1 - \frac{r_p}{r_{p\text{min}}} \right) \right] \tag{4}$$

where k is a scalar, r_p is the osculating periapsis radius, and $r_{p\text{min}}$ is near or equal to the lowest permissible value of r_p.

The scaling function is used primarily to ensure convergence to the target orbit and takes the form

$$S_{\boldsymbol{\varpi}} = \begin{cases} \left[1 + \left| \dfrac{a - a_T}{m a_T} \right|^n \right]^{\frac{1}{r}} & \text{for } \boldsymbol{\varpi} = a \\ 1 & \text{for } \boldsymbol{\varpi} = e, i, \omega, \Omega \end{cases} \tag{5}$$

where m, n, and r are scalars. The distance function is defined as

$$d(\boldsymbol{\varpi}, \boldsymbol{\varpi}_T) = \begin{cases} \boldsymbol{\varpi} - \boldsymbol{\varpi}_T & \text{for } \boldsymbol{\varpi} = a, e, i \\ \cos^{-1}\left[\cos(\boldsymbol{\varpi} - \boldsymbol{\varpi}_T)\right] & \text{for } \boldsymbol{\varpi} = \omega, \Omega \end{cases} \tag{6}$$

where the principal value, namely within $[0, \pi]$, is used for the arc cosine. The specific form of the distance function for ω and Ω is used because it provides an angular measure of the distance between two positions on a circle using the "short way round" the circle, because it is differentiable with respect to $\boldsymbol{\varpi}$ [except when $d(\boldsymbol{\varpi}, \boldsymbol{\varpi}_T) = \pi$], and because the sign of the derivative indicates whether $\boldsymbol{\varpi}$ leads or lags $\boldsymbol{\varpi}_T$ based on the short way round.

As shown above, the Q-law specifies the general form of the proximity quotient Q and the general rules for optimal thrust angles and thrust-arc locations. However, to maintain a certain degree of flexibility, the Q-law involves a set of internal parameters or weights which can be set by a mission designer to specific values. The set is composed of $\{\eta_{\text{cut}}, W_a, W_e, W_i, W_\omega, W_\Omega, W_P, m, n, r, r_{p\text{min}}, k\}$. As discussed in Sec. 4, these parameters have nominal values that should allow the Q-law to perform reasonably well for most orbit transfers. The goal of the evolutionary algorithms is to find Pareto-optimal parameter sets for any given orbit transfer problem.

3. Q-Law Optimization with Evolutionary Algorithms

Mathematically, the Q-law parameter optimization problem is expressed as

$$\text{minimize } \mathbf{y} = \{t_f(\mathbf{x}), m_p(\mathbf{x})\} \in \mathbf{Y}, \tag{7}$$

$$\text{where } \mathbf{x} = \{W_a, W_e, W_i, W_\omega, W_\Omega, W_P, m, n, r, r_{p\text{min}}, k, \eta_{\text{cut}}, \theta_i\} \in \mathbf{X}. \tag{8}$$

Here, $\mathbf{x}$ is the Q-law parameter vector, $\mathbf{y}$ the objective vector given by the required flight time (t_f) and the required propellant mass (m_p) for a given orbit transfer, $\mathbf{X}$ the decision space, and $\mathbf{Y}$ the objective space. We add one more parameter to the decision space; the initial true anomaly θ_i, which is not a Q-law parameter *per se*, but a mission-design parameter. One decision vector $\mathbf{x}_i$ leads to one candidate trajectory with a final fight time and a consumed propellant mass, that is, an objective vector $\mathbf{y}_i$. In the following paragraphs, we will describe how the optimization problem is solved with two evolutionary algorithms: a genetic algorithm and a simulated-annealing algorithm.

Genetic Algorithm

Genetic algorithms (GA), first introduced by John Holland and his colleagues [5], are search algorithms based on the mechanics of natural selection and sexual reproduction. GAs are theoretically and empirically proven to provide robust search in complex spaces. Furthermore, they are not fundamentally limited by restrictive assumptions about the search space such as continuity and existence of derivatives.

The standard GA proceeds as follows. A possible solution of a given problem is encoded as a finite string of symbols, known as the genome. An initial population of the possible solutions called individuals is generated at random or heuristically. At every evolutionary step, known as a generation, the individuals in the current population are decoded and evaluated according to some predefined quality criterion, referred to as the fitness. To form the next generation, parents are selected according to their fitness. Many selection procedures are currently in use, one of the simplest being Holland's original fitness-proportionate selection, where individuals are selected with a probability proportional to their relative fitness. This ensures that the expected number of times an individual is chosen is approximately proportional to its relative performance in the population. Thus, high-fitness individuals stand a better chance of reproducing, while low-fitness ones are more likely to disappear.

The parent selection process is followed by genetically-inspired operators to form offsprings. The most well known operators are crossover and mutation. Crossover is performed with probability p_{cross} between two selected parents, by exchanging parts of their genomes to form two offsprings; in its simplest form, substrings are exchanged after a randomly selected crossover point. This operator tends to enable the evolutionary process to move toward "promising" regions of the search space. The mutation operator is introduced to prevent premature convergence to local optima by randomly sampling new points in the search space with some probability p_{mut}. Genetic algorithms are stochastic iterative processes that are not guaranteed to converge. The termination condition may be specified as some fixed, maximal number of generations or as the attainment of an acceptable fitness level.

Nondominated Sorting

The standard GA require a ranking/evaluation scheme in the process of fitness assignment, which depends on optimization problems. Optimizing the Q-law parameters is a multi-objective optimization problem, because both propellant masses and flight times need to be minimized. In such a problem, there may not exist one solution that is best with respect to all objectives. Therefore, the goal of the multi-objective optimization problem is to determine the trade-off surface, which is a set of nondominated solution points known as Pareto-optimal or non-inferior solutions. A conventional way to solve multi-objective problems is to transform the original problem in a single-objective one, by weighting the objectives with a weight vector. However, this process tends to lead to a subgroup of Pareto-optimal solutions that is sensitive to the weight vector used in the weighting process. In contrast, the nondominated sorting process equally encourages all nondominated solutions to survive [12]. The nondominated sorting genetic algorithm (NSGA) was shown to be superior [6] to other multi-objective evolutionary algorithms such as the vector evaluated genetic algorithm (VEGA) [7], the niched Pareto genetic algorithm (NPGA) [8], and the multi-objective genetic algorithm (MOGA) [9]. Hence, we apply NSGA to optimize the Q-law parameters.

The nondominated sorting proceeds as follows. First, the nondominated individuals in the current population are identified as described in the Appendix. The same fitness value is assigned to all the individuals constituting the first nondominated front. The individuals are then ignored temporarily, and the rest of the population is processed in the same way to identify a new set of nondominated individuals constituting the second nondominated front. A fitness value that is smaller than the previous one is assigned to all the individuals of the second nondominated front. This process continues until the whole population is classified into nondominated fronts.

Simulated Annealing

Simulated annealing (SA) is a widely used and well-established optimization technique especially for high-dimensional configuration spaces [10][11]. The goal is to minimize an energy or fitness function E (in our case, the required flight time and propellant mass), which is a function of N variables (in our case, the Q-law parameters), with N being a large number. The minimization is performed by randomly changing the value of one or more of the N variables and reevaluating the energy function E. Two cases can occur: 1) the change in the variable values results in a new, lower energy function value; or 2) the energy function value is higher or unchanged. In the first scenario the new set of variable values is stored and the change accepted. In the second scenario, the new set of variable values is only stored with a certain likelihood (Boltzmann probability, including an annealing temperature). This ensures that the overall optimization algorithm will not be trapped in local minima too easily as is the case with greedy downhill optimization. As the annealing temperature decreases in the course of the optimization process, an energetically unfavorable step is less likely to be accepted (cooling schedule). The procedure is repeated until the annealing temperature has reached its end value, a preset number of iterations has been exceeded, or the energy function E has reached an acceptable level,

4. ORBIT-TRANSFER RESULTS

The parameters of the Q-law are optimized by GA and SA for five different types of orbit transfers. Table 1 lists the initial and final orbit elements, thrust characteristics, spacecraft initial mass, and central bodies associated with the five orbit

Table 1. Initial and final orbit elements, thrust characteristics, spacecraft initial masses, and central bodies associated with the orbit transfers studied in this paper. The orbit elements are given by the semimajor axis (a), the eccentricity (e), inclination (i), argument of the periapsis (ω), and longitude of the ascending node (Ω). The true anomaly (θ) is left free for both the initial and final orbit.

Case	Orbit	a (km)	e	i (degree)	ω (degree)	Ω (degree)	Thrust (N)	Specific Impulse (s)	Initial Mass (kg)	Central Body
A	Initial	7000.00	0.010	0.050	0.0	0.00	1	3100	300	Earth
	Target	42000.00	0.010	free	free	free				
B	Initial	24505.90	0.725	7.050	0.0	0.00	0.350	2000	2000	Earth
	Target	42165.00	0.001	0.050	free	free				
C	Initial	9222.70	0.200	0.573	0.0	0.00	9.3	3100	300	Earth
	Target	30000.00	0.700	free	free	free				
D	Initial	944.64	0.015	90.060	156.9	-24.60	0.045	3045	950	Vesta
	Target	401.72	0.012	90.010	free	-40.73				
E	Initial	24505.90	0.725	0.060	180.0	180.00	2	2000	2000	Earth
	Target	26500.00	0.700	116.000	270.0	180.00				

transfers termed case A, B, C, D, and E. These cases correspond to those in [4], except that for case E the plane of the initial orbit is changed by 0.12 degrees. The gravitational parameter for the orbit transfer around the Earth is set to be 398,600.5 km^3s^{-2}, while that for Vesta is 17.8 km^3s^{-2}. As is customary with the classical orbit elements, values of zero are not used for the eccentricity and inclination on account of the singularities present in Gauss's form of the variational equations. The orbit transfers range from the simpler, where few elements have target values, to the more complex, where not only do all elements have target values, but also where temporary, large sacrificial changes must be made in some elements to change more effectively other elements, until all elements converge on their target values. Recall that to effect an orbit transfer, the Q-law not only provides thrust angles but also an indication of whether to thrust or coast. Thus, the Q-law can examine the trade-off between propellant mass and flight time: To obtain short flight times, more propellant must be used, while when longer flight times are allowed, the required propellant mass is reduced. As the permitted flight time increases, eventually there are diminishing returns on the saved propellant mass, and so the flight time will typically be capped at some large-enough value for each of these transfers.

The Pareto fronts (in propellant mass and flight time) obtained with the optimized Q-law are compared with those obtained with the nominal (unoptimized) Q-control law. Furthermore, for cases A, B, C, and D, we assess how well the Pareto front of the optimized Q-law matches the performance of individual optimal transfer trajectories reported in the literature, computed using optimal control techniques (i.e. without the imposition of a feedback control law). Due to the difficulty of the optimal control problem, there is a dearth of optimal, many-revolution orbit transfers in the literature, especially when coast arcs are involved or when the transfer is complex. For each case we present the computation time needed to generate the Pareto front, and, where possible, compare this to the times needed to obtain the optimal solutions reported in the literature.

The nominal Q-law uses $W_{\mathbf{oe}} = 1$ for orbit elements with target values, $W_{\mathbf{oe}} = 0$ for orbit elements without target values, and $m = 3, n = 4, r = 2$ for the scaling function of the semimajor axis a. The penalty function to enforce minimum-periapsis-radius constraints is applied only for case D and E orbit transfers. The penalty function of the nominal Q-law uses $W_p = 1$, $k = 100$, and $r_{p\text{min}} = 300$ km for case D and $r_{p\text{min}} = 6578$ km for case E. The Pareto front of the nominal Q-law is acquired by varying the thrust effectivity threshold $\eta_{\text{cut}} \in [0, 1]$ and the initial true anomaly $\theta_i \in [0, 2\pi]$. In both the nominal Q-law and the optimized Q-law, a minimum thrust-arc length of 10 degrees is imposed, measured in true longitude $(\theta + \omega + \Omega)$.

The GA optimization uses the following GA parameters: the population size $N_p = 1000$ for case A, B, C and $N_p = 2000$ for case D and E, the number of generations $N_g = 100$, the population replacement rate $p_r = 0.1$, the crossover probability $p_c = 0.8$, the mutation probability $p_m = 0.3$. The relatively high mutation rate is chosen to preserve the diversity of the population. Each Q-law parameter is represented as a real-valued gene. The fitness of each individual is assigned according to the nondominated sorting as described in Sec. 3. Possible parents are selected by tournament (i.e., randomly pick two individuals and choose the one that is better fitted). The crossover is performed by choosing one point in the gene string at which the two strings are crossed. The mutation is performed by randomly choosing a gene in the string according to the mutation probability and resetting the gene randomly within a given range.

The SA optimization uses as its fitness function the sum of the consumed propellant mass and the flight time. The design of this fitness function results in an approximately equal optimization of both the consumed propellant mass and the flight time. Thus, a complete Pareto front cannot be expected from this fitness function. By replacing the flight time in the fitness function with the relative difference between the current flight time and a specified flight time and by varying the specified flight time, one can obtain a complete Pareto front. The SA optimization runs on a single processor, but it can be trivially parallelized by deploying N specified flight times on a cluster of N processors.

Case A Orbit Transfer

Case A is a simple coplanar, circle-to-circle orbit transfer from low Earth orbit to geostationary orbit. No periapsis constraint is imposed during the transfer, as the natural dynamics does not decrease the periapsis altitude. The maximum-permitted flight time is 500 days. Figure 1 shows the Pareto front obtained with the nominal Q-law and the optimized Q-law. Note that each solution in the Pareto front for the optimal Q-law is obtained with a different set of Q-law parameters. As shown in Fig. 1, the GA Pareto front dominates the Pareto front given by the nominal Q-law.

The Pareto-optimal solutions found by GA and SA are compared with two analytical solutions that approximately bound the problem: The Edelbaum transfer and the Hohmann transfer. The Edelbaum transfer provides an approximate lower limit for the required flight time [2], while the Hohmann transfer [13] sets an approximate lower limit for the required propellant mass. The Edelbaum transfer is a continuous-thrust, minimum-time transfer based on orbit averaging. The Hohmann transfer utilizes two thrust impulses, that is, two instantaneous large changes in velocity each without change in position. Applying thrust impulsively is much more efficient than applying it continuously over an orbit, and so the propellant required for the Hohmann transfer (assuming the thrust can be arbitrarily large) is much less than that needed for continuous thrust. In the case of low thrust, these large velocity changes can be accumulated gradually by utilizing a series of small thrust arcs. As these thrust arcs become infinitesimal in size, the propellant requirement will converge to that needed for the Hohmann transfer.

When the Q-law optimized with GA is used, the flight-time-optimal solution is about 0.04 days away from the lower limit of the flight time (14.42 days), and the propellant-optimal solution is about 0.14 kg away from the lower limit of the propellant mass (34.97 kg). In contrast, the flight-time optimal solution found by the nominal Q-law is 0.11 days away, and the propellant-optimal solution is 0.82 kg away. This comparison clearly shows that the optimization of the Q-law with GA essentially matches the theoretical flight-time and propellant bounds, having improved the Pareto front of the nominal Q-law by about 0.6% in minimum flight time and about 1.9% in minimum propellant mass. The optimized-Q-law transfer with the lowest propellant mass has a flight time of about 230 days, even though the maximum-permitted flight time is 500 days. The distance from the flight-time cap is due to the fact that the propellant mass is already very close to its minimum value, and that beyond about 250 days, the flight time becomes very sensitive to the value of $\eta_{\rm cut}$, making it difficult to populate the Pareto front beyond this flight time.

One of the limitations of the nominal Q-law for this transfer is that the nominal Q-law excludes a subgroup of Pareto-optimal solutions. As shown in Fig. 1, the nominal Q-law provides two families of Pareto-optimal solutions: one for short flight times ($14 < t_f < 17$) and the other one for long flight times ($t_f > 140$). No solutions are found for the intermediate flight times ($17 < t_f < 140$). In contrast, the GA-optimized Q-laws lead to Pareto-optimal solutions in a wide range of flight times without any significant gap. This indicates that some Q-law parameters besides $\eta_{\rm cut}$ strongly affect the trajectories to be taken.

To show which parameters are important in determining the trajectory pattern, we investigate a correlation between the optimal parameters and the flight time (or the propellant mass). While other Q-law parameters do not show much correlation, the optimal set $\{W_a, W_e, \eta_{\rm cut}\}$ show a strong correlation with the flight time, as shown in Figure 2. For example, the trajectory with flight time 50 days can be found only with $W_a = 55\%$, $W_e = 45\%$ and $\eta_{\rm cut} = 0.84$, while the rest of Q-law parameters can vary widely yet yield comparable performance. This sensitivity/correlation analysis between the Q-law parameters and the resulting trajectory suggests that the Q-law can be effectively optimized by varying only $\{W_a, W_e, \eta_{\rm cut}\}$.

We present examples of the three families of trajectories in the Pareto front: for flight-time optimal solutions, propellant-optimal solutions, and intermediate-flight-time solutions found by the optimized Q-law. The flight-time-optimal trajectory is roughly a circular spiral, increasing the semimajor axis while maintaining the eccentricity close to zero, as shown in Figs. 3 and 6. The propellant-optimal trajectory takes quite a different form, maintaining the same periapsis until the apoapsis radius slightly surpasses the target value, and then increasing the periapsis radius to near the target value, before finally driving both periapsis and apoapsis radii to their target values, as shown in Figs. 4 and 6. As expected, the intermediate-flight-time trajectory, shown in Figs. 5 and 6, is a hybrid between the flight-time optimal and the propellant-optimal trajectories, increasing the periapsis and apoapsis simultaneously. The trajectory initially increases both the eccentricity and the semimajor axis, and later reduces the eccentricity while continuing to increase the semimajor axis.

Case B Orbit Transfer

Case B is a transfer from a slightly-inclined geostationary transfer orbit to geostationary orbit. The maximum-permitted

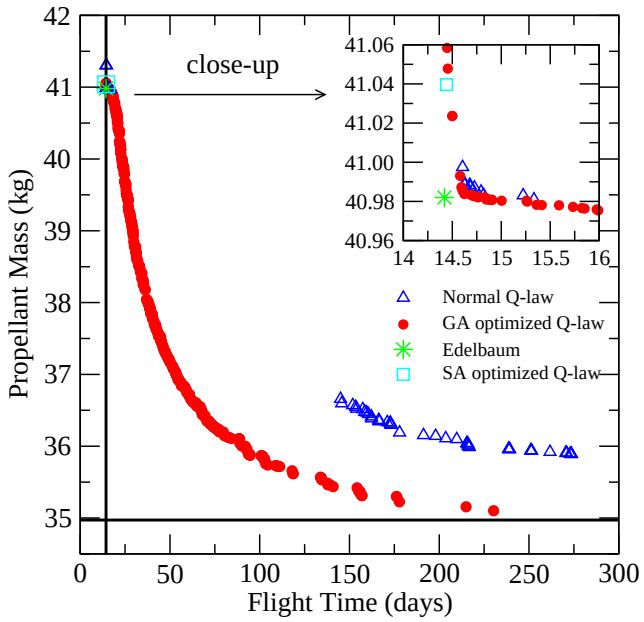

Figure 1. Case A: Trade-off between propellant mass and flight time. The Pareto fronts generated by the nominal Q-law and the GA optimized Q-law are plotted in comparison with the lower bounds (solid lines) of the required flight time and propellant mass given by the Edelbaum and Hohman transfers, respectively. A flight-time optimal solution found by the Q-law optimized by SA is also plotted.

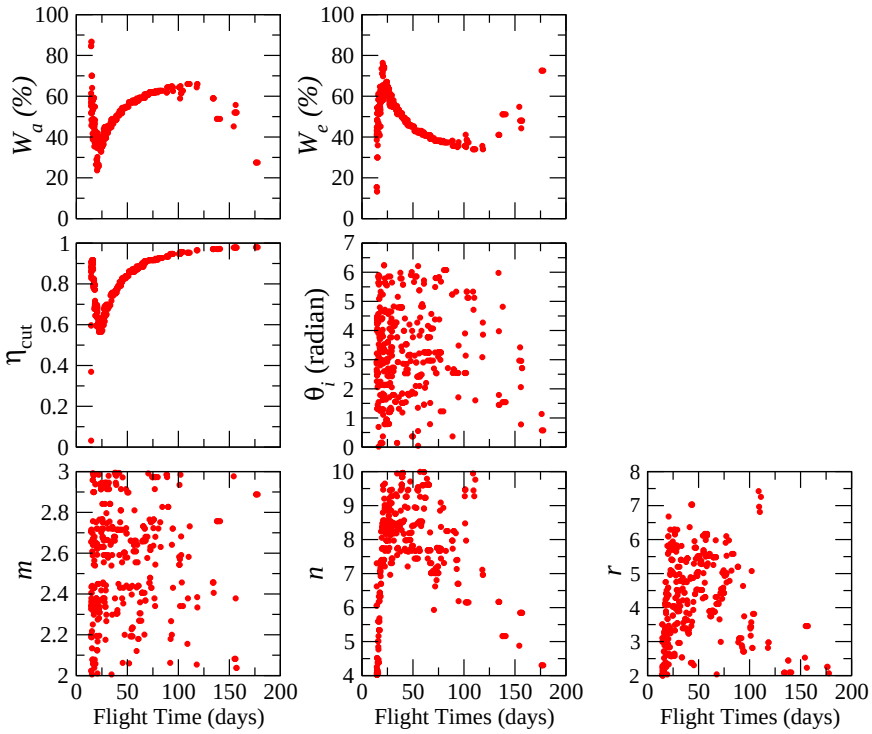

Figure 2. Case A: Optimal Q-law parameters found by GA with respect to flight time. There is a strong correlation between $\{W_a, W_e, \eta_{\rm cut}\}$ and the flight time, while other Q-law parameters show a weak correlation.

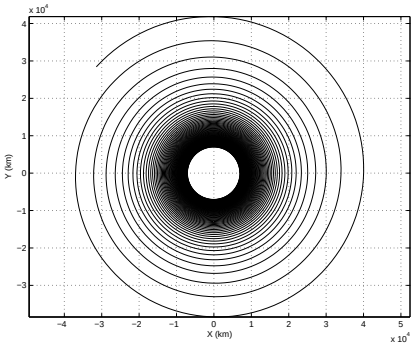

Figure 3. Case A: flight-time-optimal trajectory with flight time 14.5 days.

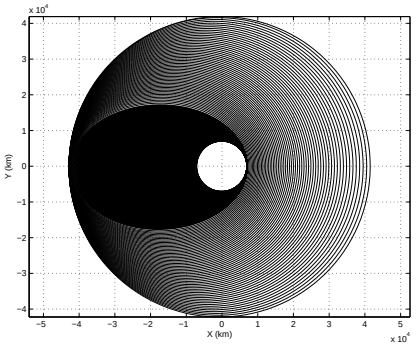

Figure 4. Case A: propellant-optimal trajectory with flight time 230 days.

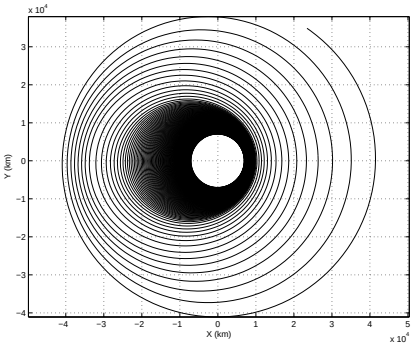

Figure 5. Case A: intermediate-flight-time trajectory with flight time 30 days.

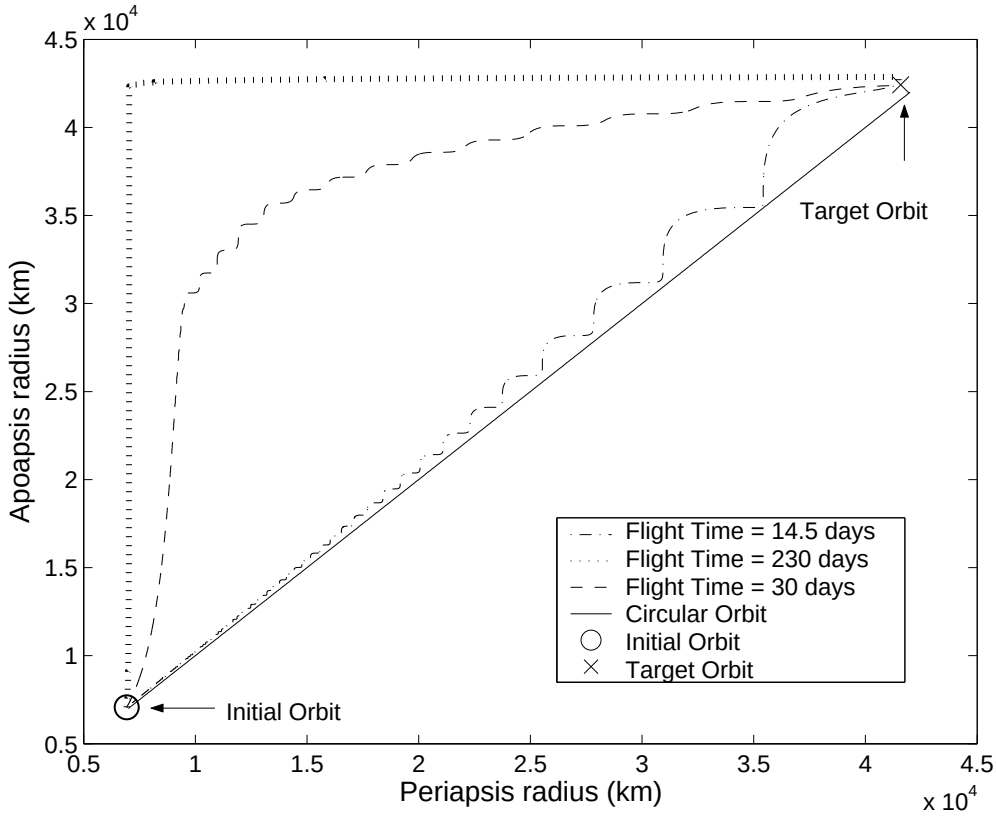

Figure 6. Case A: Time history of apoapsis and periapsis radii for three trajectories with flight times of 14.5 days (dashed dot line), 230 days (dotted line), and 30 days (dashed line).

flight time is 1500 days. Figure 7 shows the trade-off between propellant-mass and flight-time for this transfer. In comparison with the Pareto front generated with the nominal Q-law, the improvement of the Pareto front with the optimized Q-law is dramatic. A propellant savings of about 5-15% is achieved with the optimized Q-law. To verify the quality of the improved Pareto front, we compare it with the two optimal trajectories found by Geffroy and Epenoy using an orbit averaging technique [14]. The inset of Fig. 7 shows that our Pareto-optimal solutions are as good as the solutions found by Geffroy and Epenoy.

An analysis of the correlation between the optimal Q-law parameters and the flight time is shown in Figure 8. The dense populations of optimal W_a around 10%, optimal W_e around 20%, and W_i around 70% show that the nominal Q-law ($W_a = W_e = W_i$) is not an optimal choice. As expected, the thrust effectivity threshold $\eta_{\rm cut}$ is the important parameter to control the flight time. Other Q-law parameters (m, n, r) and the initial true anomaly (θ_i) show a weak correlation with the flight time, indicating that these parameters are not as critical as W_a, W_e, W_i, and $\eta_{\rm cut}$ in the Q-law optimization.

Case C Orbit Transfer

Case C is a transfer from a low-eccentricity elliptic orbit to a coplanar, high-eccentricity, larger elliptic orbit, with a maximum-permitted flight time of 20 days. Figure 9 shows the trade-off between propellant mass and flight time for this transfer. The Pareto front for the nominal Q-law is obtained by varying the thrust effectivity threshold $\eta_{\rm cut} \in [0, 1]$ and the initial true anomaly $\theta_i \in [0, 2\pi]$. The Pareto front for the GA optimized Q-law is generated by optimizing $\{W_a, W_e, m, n, r, \eta_{\rm cut}, \theta_i\}$. The GA optimized Q-law provides a better estimation of the Pareto front than the nominal Q-law particularly for short flight times. For longer flight times, the Pareto front of the nominal Q-law is truncated at a flight time of about 5.3 days due to the minimum thrust-arc length constraint — when this constraint is removed, the Pareto front of the nominal Q-law is improved and closely follows the optimized Q-law front. Several solutions found with the optimization tool Mystic are plotted for comparison. Mystic uses the static/dynamic control algorithm [15] [16]. The comparison shows that the Pareto front generated by the optimized Q-law is as good as the Mystic solutions.

The optimal Q-law parameters found by GA are plotted with respect to flight time in Fig. 10. The optimal W_a, W_e and $\eta_{\rm cut}$ are strongly correlated to the flight time, while other Q-law parameters show a weak correlation. In general, flight-time-optimal solutions have $W_e/W_a > 1$, while propellant-optimal solutions have $W_e/W_a < 1$. This means that the flight-time-optimal solutions emphasize the eccentricity target while the propellant-optimal solutions emphasize the semi-major axis target.

Case D Orbit Transfer

Case D is roughly a circle-to-circle orbit transfer around the asteroid Vesta, involving a small plane change. The flight time is capped at 300 days. Figure 11 shows the trade-off between propellant mass and flight time for this transfer. The Pareto front of the nominal Q-law is obtained by varying the thrust effectivity threshold $\eta_{\rm cut} \in [0, 1]$ and the initial true anomaly $\theta_i \in [0, 2\pi]$. The Pareto fronts of the GA optimized Q-law are generated in three different ways: the first Pareto front (GA Q-law I) is obtained by optimizing $\{W_a, W_e, W_i, W_\Omega, \eta_{\rm cut}, \theta_i\}$, the second Pareto front (GA Q-law II) by optimizing $\{W_a, W_e, W_i, W_\Omega, \eta_{\rm cut}, \theta_i, m, n, r\}$, and the third Pareto front (GA Q-law III) by optimizing $\{W_a, W_e, W_i, \eta_{\rm cut}, \theta_i, m, n, r, W_P, r_{p\rm min}, k\}$. In comparison with the nominal Q-law, the GA optimized Q-law improves an estimation of the Pareto front for all the flight times considered. The GA optimized Q-law leads to a propellant mass savings as large as 16%. More promisingly, the Pareto-optimal solutions found with the optimized Q-law are as good as the solution found by Whiffen using the static/dynamic control algorithms coded in Mystic [16] [17].

Among the three GA optimization schemes described above, GA Q-law II and GA Q-law III outperform GA Q-law I but the difference between GA Q-law II and GA Q-law III is insignificant. This result indicates that the trajectory does not depend strongly on $\{W_P, r_{p\rm min}, k\}$ (the parameters of the penalty function for the minimum periapsis constraint) and thus an accurate Pareto front can be obtained by optimizing only $\{W_a, W_e, W_i, \eta_{\rm cut}, \theta_i, m, n, r\}$. The difference between the Pareto fronts generated by GA Q-law I and GA Q-law II (or III) becomes smaller as the flight time becomes longer. This sheds some light on the effect of the Q-law parameters $\{m, n, r\}$ on the Q-law performance. The parameters $\{m, n, r\}$ are introduced for the scaling function in the semimajor axis to ensure the convergence of transfers which involve an increase in the semimajor axis. However, the semimajor axis steadily decreases in this orbit transfer, suggesting that the scaling function is not needed. Therefore, it is preferable to select a parameter set $\{m, n, r\}$ that yields the smallest possible modification to the distance function.

The optimal Q-law parameters found with GA are plotted with respect to the flight time in Figure 12. Optimal W_a, W_e, W_i, W_Ω are normalized to make the sum to be 100%. The Q-law optimization shows a greater correlation for $\{W_a, W_e, W_i, W_\Omega, \eta_{\rm cut}, \theta_i, m, n, r\}$ than for $\{W_p, r_{p\rm min}, k\}$. This explains the similarity between the Pareto front generated with GA Q-law II and the Pareto front generated by GA Q-law III. As in other transfers, this transfer shows a strong correlation between $\eta_{\rm cut}$ and the flight time. However, the correlation does not follow the monotonous trend that a larger $\eta_{\rm cut}$ leads to a longer flight time. The optimal $\eta_{\rm cut}$ shows a discontinuity around a flight time 60 days. The discontinuity also appears in other optimal Q-law parameters such as W_a, W_e, W_Ω. This indicates that the pattern of the trajectory changes around this flight time.

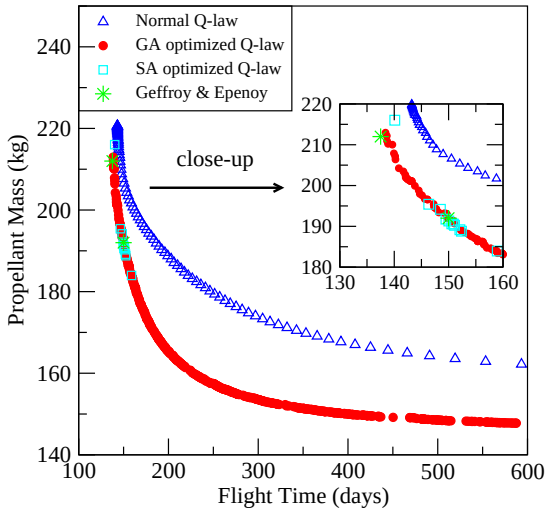

Figure 7. Case B: Trade-off between propellant mass and flight time. The Pareto fronts generated by the nominal Q-law and the GA/SA optimized Q-law are plotted in comparison with the Pareto-optimal solutions found by

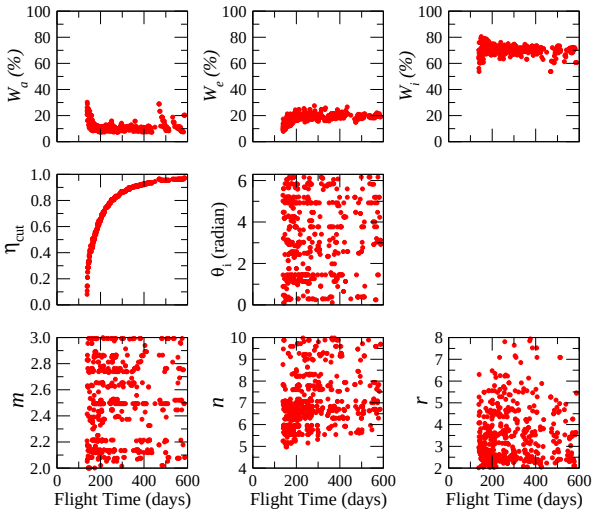

Figure 8. Case B: Optimal Q-law parameters found by GA with respect to flight time. Optimal W_a, W_e, W_i are normalized to make their sum to be 100%.

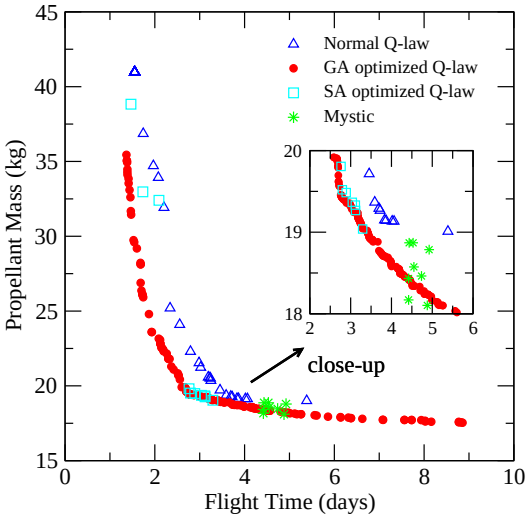

Figure 9. Case C: Trade-off between propellant mass and flight time. The Pareto fronts generated by the nominal Q-law and the GA/SA optimized Q-law are plotted in comparison with several Pareto-optimal solutions found with the optimization tool Mystic.

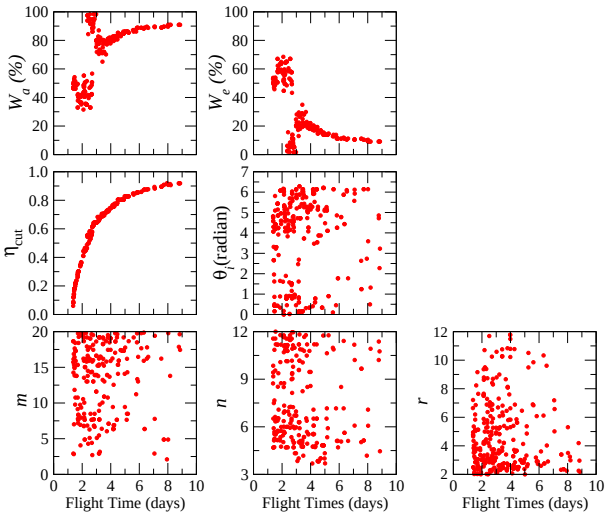

Figure 10. Case C: Optimal Q-law parameters found by GA with respect to flight time. A strong correlation between $\{W_a, W_e, \eta_{\rm cut}\}$ and the flight time is observed, while other Q-law parameters show a weak correlation.

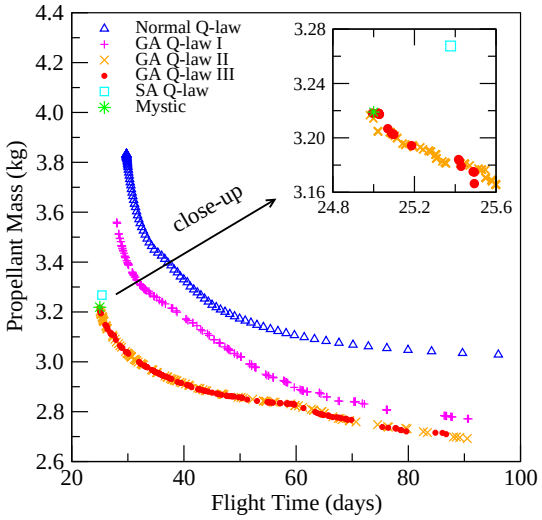

Figure 11. Case D: Trade-off between propellant mass and flight time. The Pareto fronts are obtained with the nominal Q-law and with the Q-law optimized with GA. A flight-time optimal solution is found by the Q-law optimized with SA. A Pareto-optimal solution found by Mystic is also plotted for comparison.

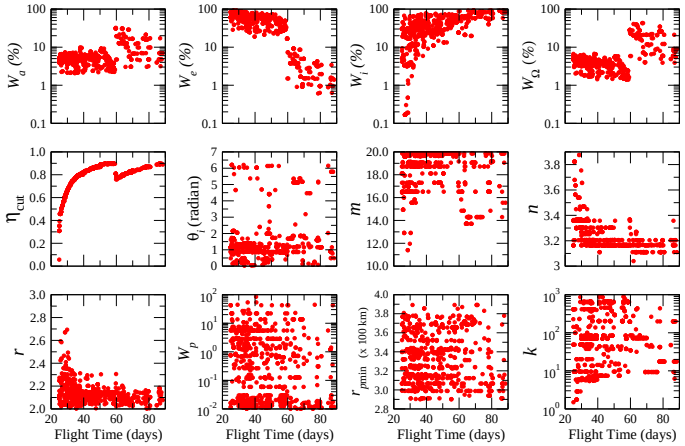

Figure 12. Case D: Optimal Q-law parameters found with GA with respect to the flight time. The overall distribution of the optimal parameters shows that the Q-law performance is more sensitive to the choice of $\{W_a, W_e, W_i, W_\Omega, \eta_{\rm cut}, \theta_i, m, n, r\}$ than $\{W_P, r_{p\rm min}, k\}$.

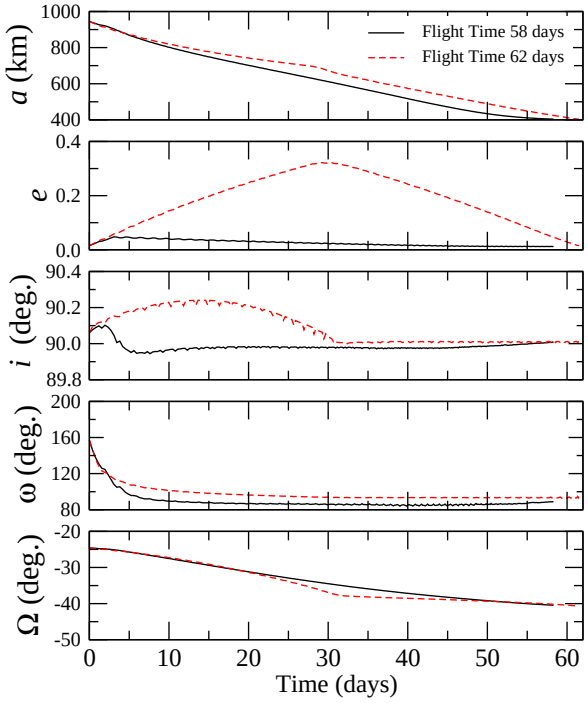

Figure 13. Case D: Orbit elements as a function of time for a Pareto-optimal trajectory with flight times 58 days (just below the discontinuity point of the optimal $\eta_{\rm cut}$ shown in Fig. 12) and 62 days (just above the discontinuity point). A large difference in the time history of the eccentricity between the two trajectories is observed, while other orbit elements show little difference.

To understand the cause of the discontinuity of the optimal Q-law parameters, we examine the trajectory for a flight time just below the discontinuity point (T1) and that for a flight time just above the discontinuity point (T2). Figure 13 shows orbit elements as a function of time during the orbit transfer. The two trajectories show a significant difference in the time history of the eccentricity, while other orbit elements (a, i, ω, Ω) show a small difference. T1 keeps the eccentricity close to zero all time, but T2 shows a large increase and decrease of the eccentricity during the orbit transfer. This trend is similar to that observed in Case A, where the circular spiral trajectory (Edelbaum-type transfer) is flight-time optimal and the elliptic trajectory (Hohmann-type transfer) is propellant optimal. The two types of trajectories can be obtained with the Q-law by either emphasizing the eccentricity target or not. This result is also observed in the distribution of the optimal W_e in Figure 12. The optimal W_e is greater for short-flight-time solutions than for long-flight-time solutions.

Case E Orbit Transfer

Case E is a transfer from a geostationary transfer orbit to a retrograde, Molniya-type orbit, involving a large plane change. The maximum-permitted flight time is 300 days. Figure 14 shows the trade-off between propellant mass and flight time for this transfer. The Pareto front for the nominal Q-law is obtained with varying $\eta_{\rm cut} \in [0, 1]$ and the initial true anomaly $\theta_i \in [0, 2\pi]$. Three Pareto fronts are generated with GA optimization as follows: the first Pareto front (GA-Q-law I) by optimizing $\{W_a, W_e, W_i, W_\omega, W_\Omega\}$. the second Pareto front (GA Q-law II) by optimizing $\{W_a, W_e, W_i, W_\omega, W_\Omega, m, n, r, \eta_{\rm cut}, \theta_i\}$, and the third Pareto front (GA Q-law III) by optimizing $\{W_a, W_e, W_i, W_\omega, W_\Omega, m, n, r, \eta_{\rm cut}, \theta_i, W_P, r_{p\rm min}, k\}$. The GA optimized Q-law provides a better estimation of the Pareto front than the nominal Q-law for all the flight times considered. A propellant mass savings as large as 30% is obtained with the GA optimized Q-law. Like Case D, GA Q-law II and GA Q-law III outperform GA Q-law I in this case, while the difference between GA Q-law II and III is insignificant. This result reflects the degree of influence of each Q-law parameter on the Q-law performance. The difference between GA Q-law I and GA Q-law II (or III) becomes larger as the flight time increases in contrast to Case D.

The optimal Q-law parameters found with GA are plotted with respect to the flight time in Figure 15. The overall distribution of the optimal Q-law parameters shows the greater sensitivity of the Q-law performance to $\{W_a, W_e, W_i, W_\Omega, \eta_{\rm cut}\}$ than to $\{m, n, r, W_p, r_{p\rm min}, k\}$. The optimal $\eta_{\rm cut}$ shows a strong correlation with flight time as was found for other transfers. A strong preference for the relative size hierarchy $W_i > W_\Omega > W_a > W_\omega > W_e$ is observed for all flight times.

Case E specifies changes in all orbit elements, making it the most complicated transfer among the five transfers studied here. We examine how the change of each orbit element interacts with other orbit-element changes. Figure 16 shows the time history of each orbit element for four different Pareto-optimal trajectories found by GA Q-law III. For the all four trajectories, the plane changes (i.e. i, ω, Ω) occur when the semimajor axis nearly reaches the maximum values, and the increase of the semimajor axis is accompanied by an increase of the eccentricity. This behavior stems from the orbit-transfer energetics, in which the larger apoapsis radius (i.e. larger semimajor axis and larger eccentricity) reduces the cost of the plane change in terms of propellant consumption.

Figure 16 also unveils a general trend in orbit-element changes with respect to the flight time. The trajectory with a longer flight time involves a larger change of the semimajor axis and a later start of the plane change. For example, the shortest-flight-time trajectory (the solid line) exhibits an early start of the plane change as the semimajor axis peaks at 50,000 km. In contrast, the longest-flight-time trajectory (the line with circles) shows almost no plane change until the semimajor axis reaches its maximum 100,000 km. The difference is directly related to the orbit-transfer energetics, in which the plane change with a larger apoapsis radius is propellant efficient. The longer flight-time trajectory takes better advantage of the energetics. The top panel of Fig. 16 illustrates the time history of propellant usage during the transfer. The shortest-flight-time trajectory uses propellant with an almost constant rate. The longer-flight-time trajectories use propellant with a lower rate during the first stage of the semimajor-axis increase followed by a higher rate of propellant consumption in the second stage of the plane change.

Computational Requirement

The computation time required to obtain the Pareto front for each orbit transfer is listed in Table 2. The computer used for the GA calculation is a 32 node Beowulf cluster with 3.06 GHz Pentium 4 microprocessor, while the computer used for the SA calculation is a 31 node Beowulf cluster with 800 MHz Pentium III microprocessor. In the GA calculation, Case C requires a relatively short computation time because the evaluation of each Q-law takes less time due to the short flight time in this orbit transfer. Beside Case C, the required computation time is between 11.8 to 41.3 CPU hours. For Case A, B, and C, the GA computation evaluates 10,000 sets of Q-law parameters, while for Case D and E it evaluates 20,000 sets of Q-law parameters. Therefore, the time to evaluate one set of Q-law parameters (equivalently to obtain a candidate trajectory and to assign its fitness) is only about 0.001 CPU hours (0.1 minutes) on average.

In addition to the efficient evaluation of candidate Q-laws/trajectories, GA and SA are amenable to a parallel computing implementation thanks to the independent evaluation of each candidate Q-law/trajectory in the population/ensemble. The parallel computation significantly reduces the wall-clock time for a given computational load. For this work, the GA computation was performed on 10 processors in parallel, thus requiring a wall-clock time that is one

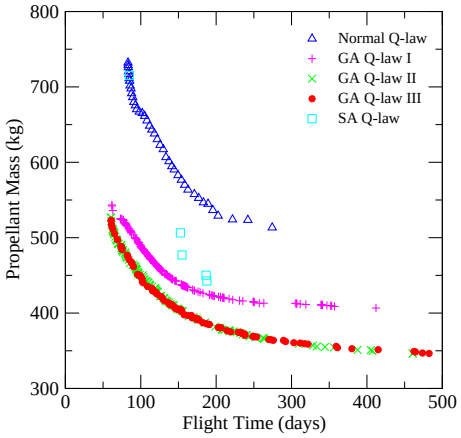

Figure 14. Case E: Trade-off between propellant mass and flight time. The Pareto fronts are obtained with the nominal Q-law and the Q-law optimized with GA and SA.

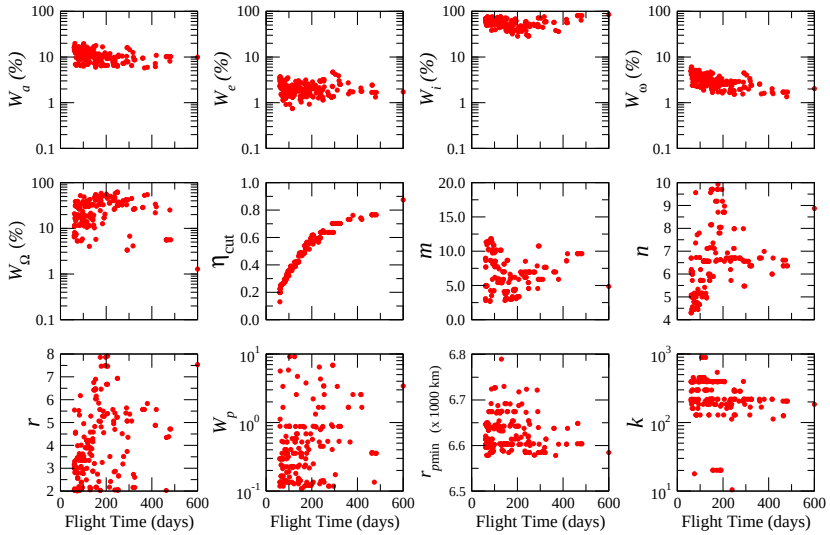

Figure 15. Case E: Optimal Q-law parameters found by GA with respect to flight time. There is a strong correlation between $\eta_{\rm cut}$ and the flight time, while other Q-law parameters show a weak correlation. A strong preference for the relative size hierarchy $W_i > W_\Omega > W_\omega > W_e$ is observed for all flight times.

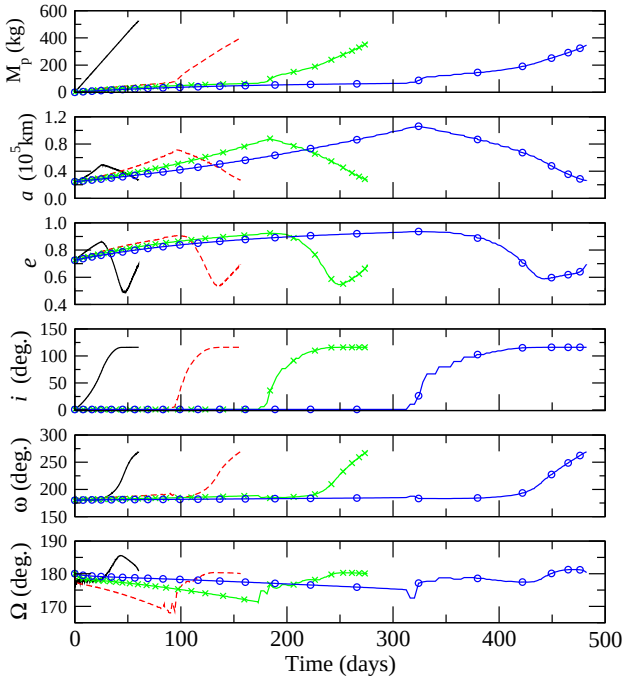

Figure 16. Case E: Consumed propellant mass and orbit elements as a function of time for four Pareto-optimal trajectories among the solutions found by GA Q-law III. The solid line is the trajectory with flight time 60 days, the dashed line is the trajectory with flight time 156 days, the line with x symbols is the trajectory with flight time 275 days, and the line with circles is the trajectory with flight time 482 days. As a general pattern, the trajectory with a longer flight time involves a larger change of the semimajor axis (a) and a later change of the inclination (i) and the argument of the periapsis (ω).

Table 2. Computation times required to obtain a Pareto front with the Q-law optimized with GA and SA for each orbit transfer. SA computation was performed in a single processor, while GA computation was performed on ten processors in parallel and thus required wall-clock time that is one tenth the listed computation time.

Orbit Transfer Case	Computation Time (CPU hours) GA	SA
A	11.8	5.2
B	13.3	22.5
C	1.0	44.3
D	25.8	68.9
E	41.3	38.0

tenth of the computation time listed in Table 2. It is the short wall-clock time (1.2 – 4.1 hours) that makes our optimization method attractive as a guiding tool for the early stage of mission design where many possible scenarios need to be evaluated. It is important to note that our method produces a Pareto front (i.e., a group of Pareto-optimal solutions) within a few hours, while other optimization algorithms tend to require a similar amount of computational time as well as some user interaction to acquire just a single Pareto-optimal trajectory: The Mystic solutions of Case C each typically took between 6 and 24 hours to run (although one took about a week), and the Mystic solution of Case D took about a half day [4] [16].

5. CONCLUSIONS

For the design and optimization of trajectories powered by low-thrust propulsion, we have developed an efficacious and efficient method to obtain approximate propellant and flight-time requirements and Pareto-optimal trajectories. The method involves a two-level optimization process: i) Lyapunov-optimal thrust angles and locations are determined with the Q-law, ii) the Q-law is optimized with two evolutionary algorithms: a genetic algorithm and a simulated-annealing-related algorithm. We have applied our method to four different types of orbit transfers around the Earth and one orbit transfer around the asteroid Vesta. The optimization of the Q-law yields the greatest benefit in the case of the most complex of the five orbit transfers considered, although less complex cases also benefit. The resulting Pareto front with the optimized Q-law shows a propellant savings as large as 30% in comparison with the nominal Q-law, and the Pareto front contains the optimal solutions found by other trajectory optimization algorithms.

In optimization problems, there is always a trade-off between the optimization quality and the computational requirement. Most of the efficient/fast optimization tools tend to yield low-quality solutions while high-quality optimization tools tend to require large computational resources. Both high quality of optimization and low computational requirement are needed in the early stages of mission design, where many possible scenarios are considered. Our method offers both the high optimization quality and the high computational efficiency. The trajectory quality of our method is shown to be as good as that of other state-of-the-art optimization tools. Our method yields not only a few Pareto-optimal trajectories but also an accurate Pareto front for a given orbit transfer within a few hours of computation time. The computational efficiency arises from both the efficiency of the Q-law in obtaining a candidate trajectory and the natural parallelism of GA/SA computation in evaluating a population/ensemble of candidate Q-laws/trajectories.

6. ACKNOWLEDGMENTS

The authors thank Christoph Adami, Van Dang, and Didier Keymeulen for useful discussions. This work was performed at the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration. The research was supported by JPL's R&TD program. This work is a part of the large effort by the JPL's Evolvable Computation Group to develop evolutionary computational techniques to design and optimize complex space systems and thus to improve on human-design of space systems [18].

7. APPENDIX

Mathematically, a multi-objective optimization problem is expressed as

$$\text{minimize } \mathbf{y} = \{y_1(\mathbf{x}), \cdots, y_M(\mathbf{x})\} \in \mathbf{Y}, \quad \text{(A.1)}$$

$$\text{where } \mathbf{x} = \{x_1, \cdots, x_N\} \in \mathbf{X}, \quad \text{(A.2)}$$

and $\mathbf{x}$ is the N dimensional decision vector, $\mathbf{y}$ the M dimensional objective vector, $\mathbf{X}$ the decision space, and $\mathbf{Y}$ the objective space.

Within the multi-objective optimized problem, a nondominated solution is the solution that is not dominated by any other feasible solutions. The condition for the solution $\mathbf{x}^a$ to dominate $\mathbf{x}^b$ is given by [6] [12],

$$\begin{aligned} & \forall\, i \in \{1, \cdots, M\}, y_i(\mathbf{x}^a) \leq y_i(\mathbf{x}^b) \\ \wedge \quad & \exists\, i \in \{1, \cdots, M\}, y_i(\mathbf{x}^a) < y_i(\mathbf{x}^b). \end{aligned} \quad \text{(A.3)}$$

The second condition ensures that $\mathbf{y}(\mathbf{x}^a) \neq \mathbf{y}(\mathbf{x}^b)$.

REFERENCES

[1] D.F. Lawden "Optimal Programming of Rocket Thrust Direction," *Astronautica Acta*, Vol.1, 1955, pp.41–56.

[2] T.N. Edelbaum, "Optimum Power-Limited Orbit Transfer in Strong Gravity Fields," *AIAA Journal*, **3**, 921–925, 1965.

[3] A.E. Petropoulos, "Simple Control Laws for Low-Thrust Orbit Transfers," *AAS/AIAA Astrodynamics Specialist Conference*, AAS Paper 03-630, 2003.

[4] A.E. Petropoulos, "Low-Thrust Orbit Transfers Using Candidate Lyapunov Functions with a Mechanism for Coasting," *AIAA/AAS Astrodynamics Specialist Conference*, AIAA Paper 2004-5089, 2004.

[5] J.H. Holland, *Adaptation in Natural and Artificial Systems*, The University of Michigan Press, Ann Arbor, Michigan, 1975.

[6] A.H.F. Dias and J.A. de Vasconcelos, "Multiobjective Genetic Algorithms Applied to Solve Optimization Problems," *IEEE Trans. Magn.*, **38**, 1133–1136, 2002.

[7] J.D. Schaffer, "Multiple Objective Optimization with Vector Evaluated Genetic Algorithms," Ph.D. thesis, Vanderbilt University, 1984.

[8] J. Horn, N. Nafpliotis, and D.E. Goldberg, " A niched pareto genetic algorithm for multiobjective optimization," *Proc. 1st IEEE Conf. Evolutionary Computation*, **1**, 82–87,1994.

[9] C.M. Fonseca and P.J. Fleming, "Genetic algorithms for multiobjective optimization", *Proc. 5th Int. Conf. Genetic Algorithms*, 416-423, 1993.

[10] Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E, "Equation of State Calculation by Fast Computing Machines," *J. of Chem. Phys.*, **21**, 1087–1091, 1953.

[11] Kirkpatrick, S., Gelat, C.D., Vecchi, M.P., "Optimization by Simulated Annealing," *Science*, **220**, 671–680, 1983.

[12] N. Srinivas and K. Deb, "Multiobjective optimization using nondominated sorting in genetic algorithms," *Evolutionary Computation*, **2**, 221–248, 1994.

[13] R.H. Battin, *An Introduction to the Mathematics and Methods of Astrodynamics*, 1st ed. 4th printing, AIAA, New York, 1987, pp.488–489.

[14] S. Geffroy and R. Epenoy, "Optimal Low-Thrust Transfers with Constraints - Generalization of Averaging Techniques," *Astronautica Acta*, **41**, 133–149, 1997.

[15] G.J. Whiffen and J. A. Sims, "Application of a Novel Optimal Control Algorithm to Low-Thrust Trajectory Optimization," *AAS/AIAA Space Flight Mechanics Meeting*, AAS Paper 01-209, 2001.

[16] G.J. Whiffen, "Optimal Low-Thrust Orbit Transfers around a Rotating Non-Spherical Body" *AAS/AIAA Space Flight Mechanics Meeting*, AAS Paper 04-264, 2004.

[17] In Ref. [16], solutions are sought to a different orbit transfer around Vesta than the one sought here. Although the initial orbits are identical, the final orbits differ. Ref. [16] shows that the transfer being solved there is infeasible in 25 days when Vesta is treated as a point mass (but not when a full gravity field is included); the final orbit used here in Case D is the orbit reached in Ref. [16] after 25 days of thrusting. This 25-day transfer of Ref. [16] is an optimal transfer between the initial orbit and the orbit reached after the 25 days, and so we compare our Case D to it.

[18] R.J. Terrile, C. Adami, S.N. Chao, V.T. Dang, M.I. Ferguson, W. Fink, T.L. Huntsberger, G. Klimeck, M.A. Kordon, S. Lee, P. von Allmen, and J. Xu "Evolutionary Computation Technologies for Space Systems," *IEEE Aerospace Conference Proceedings*, March 2005.

BIOGRAPHY

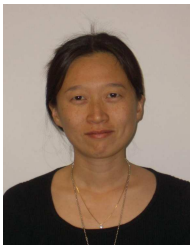

Seungwon Lee *is a member of technical staff in the Applied Cluster Computing Technologies Group at the Jet Propulsion Laboratory. Her research interest includes genetic algorithms, low-thrust trajectory design, nanoelectronics, quantum computation, materials simulation, parallel cluster computing, and advanced scientific software modernization techniques. Seungwon received her B.S. and M.S. in Physics from the Seoul National University in 1995 and 1997, and her Ph.D. in Physics from the Ohio State University in 2002. Her dissertation focused on the study of the electro-optical properties of semiconductor nanostructures. Her work is documented in numerous journals and conference proceedings.*

Paul von Allmen *is the supervisor of the Applied Cluster Computing Technologies Group and a senior researcher at the Jet Propulsion Laboratory. Paul is currently leading research in quantum computing, thermoelectric and nonlinear optics material and device design, nano-scale chemical sensors, semiconductor optical detectors, and low-thrust trajectory optimization. Prior to joining JPL in 2002, Paul worked at the IBM Zurich Research Lab on semiconductor diode lasers, at the University of Illinois on the silicon device sintering process, and at the Motorola Flat Panel Display Division as manager of the theory and simulation group and as principal scientist on micro-scale gas discharge UV lasers. Paul received his Ph.D. in physics from the Swiss Federal Institute of Technology in Lausanne, Switzerland in 1990. His work is documented in numerous publications and patents.*

Wolfgang Fink *is a Senior Researcher at NASA's Jet Propulsion Laboratory, Pasadena, CA, Visiting Research Assistant Professor of both Ophthalmology and Neurological Surgery at the University of Southern California, Los Angeles, CA, and Visiting Associate in Physics at the California Institute of Technology, Pasadena, CA. His research interests include theoretical and applied physics, biomedicine, astrobiology, computational field geology, and autonomous planetary and space exploration. Dr. Fink obtained an M.S. degree in Physics from the University of Göttingen in 1993, and a Ph.D. in Theoretical Physics from the University of Tübingen in 1997. His work is documented in numerous publications and patents.*

Anastassios E. Petropoulos *is a Senior Member of the Engineering Staff in the Outer Planet Mission Analysis Group at the Jet Propulsion Laboratory. He is actively involved in developing algorithms for designing low-thrust orbit transfers, and also in single-body and multi-body low-thrust mission design. In 1991 he received his B.S. in Aeronautical and Astronautical Engineering from the Massachusetts Institute of Technology, and in 1993 and 2001 he received his M.S. and Ph.D. degrees from Purdue University. His dissertation focused on the problem of low-thrust, gravity-assist trajectory design.*

Richard J. Terrile *created and leads the Evolutionary Computation Group at the Jet Propulsion Laboratory. His group has developed genetic-algorithm based tools to improve on human design of space systems and has demonstrated that computer-aided design tools can also be used for automatedinnovation and design of complex systems. He is an astronomer, the Mars Sample Return Study Scientist, the JIMO Deputy Project Scientist and the co-discoverer of the Beta Pictoris circumstellar disk. Dr. Terrile has B.S. degrees in Physics and Astronomy from the State University of New York at Stony Brook and an M.S. and a Ph.D. in Planetary Science from the California Institute of Technology in 1978.*